

Media Mix Model Calibration With Bayesian Priors

Yingxiang Zhang, Mike Wurm, Eddie Li, Alexander Wakim, Joseph Kelly,
Brenda Price, Ying Liu

Google LLC.

November 2023

Abstract

Effective model calibration is a critical and indispensable component in developing Media Mix Models (MMMs). One advantage of Bayesian-based MMMs lies in their capacity to accommodate the information from experiment results and the modelers' domain knowledge about the ad effectiveness by setting priors for the model parameters. However, it remains ambiguous about how and which Bayesian priors should be tuned for calibration purpose. In this paper, we propose a new calibration method through model reparameterization. The reparameterized model includes Return on Ads Spend (ROAS) as a model parameter, enabling straightforward adjustment of its prior distribution to align with either experiment results or the modeler's prior knowledge. The proposed method also helps address several key challenges regarding combining MMMs and incrementality experiments. We use simulations to demonstrate that our approach can significantly reduce the bias and uncertainty in the resultant posterior ROAS estimates.

1 Introduction

Media Mix Model (MMM) is a top-level modeling tool that utilizes statistics and historical data to understand what drives sales. It measures media investment efficiency on top of baseline sales and other external factors that affect sales (e.g. seasonality, pricing, economy, etc.). MMMs generally require modeling with causal inference assumptions and historical data to generate modeled results for incremental sales, and may fall short in precisely measuring effectiveness, especially for smaller media channels and channels with little spend variation (Chan and Perry [2017]). Incrementality experiments, on the other hand, leverage randomized controlled experiments to compare the change in consumer behavior between groups that are exposed or withheld from marketing activity while keeping all other factors constant. This presents a way of rigorously developing ground truth data on causality and provides a rigorous view of the incremental value brought by the marketing investment. In contrast to MMMs, which measure the effectiveness of all media channels over a longer period of time, incrementality experiments provide results for incremental sales at a specific point in time and are generally designed to measure the effectiveness of a specific media channel.

Modern MMMs often employ Bayesian modeling methodologies, for example Geo-level Bayesian Hierarchical Media Mix Modeling (GBHMMM, Sun et al. [2017]). One advantage of Bayesian MMMs lies in their capacity to accommodate modelers’ prior domain knowledge or beliefs about the effectiveness of various media channels. This can be particularly valuable when dealing with limited or noisy data for specific channels. Incrementality experiments provide excellent opportunities to use priors and enhance the model’s accuracy. By leveraging these experiments as priors, one can ensure a closer alignment between the model’s outputs and the actual incremental value of the channels.

However, the process of integrating incrementality experiment results, particularly Return on Ads Spend (ROAS), into MMMs via prior adjustment remains ambiguous. A conventional approach is the manual calibration method, which is an iterative process that involves trying different combinations of prior distributions for model parameters such that the implied prior distribution of ROAS aligns with experiments. This method can be time-consuming and may require significant expertise (Runge et al. [2023]). Moreover, questions remain as to how to calibrate MMMs when the short-term media effects measured by experiments are not representative of the long-term ads effectiveness measured by MMMs, and how to calibrate MMMs when multiple experiments are conducted for the same channel.

In this paper, we propose a new calibration method with Bayesian priors through reparameterization, which is developed based on GBHMMM (Sun et al. [2017]) and can be generalized to other Bayesian-based MMMs. The proposed calibration method provides a straightforward way to incorporate the prior domain knowledge about the effectiveness of media channels and the incrementality experiment results into MMMs. It directly addresses key advertiser queries: How to integrate prior knowledge or experiment results into MMMs, how to calibrate MMMs when the experiment results do not represent the modeling-window ROAS that MMMs aim to measure, and how to perform calibration with multiple experiments. Lastly, the calibration method can potentially enhance the accuracy and reduces the bias and uncertainty in MMM’s ROAS estimates, as demonstrated in our simulation study.

It is worth noting that incrementality experiments are not the only source of information for MMM calibration. Priors can encode various types of prior beliefs, such as ROAS estimates from industry benchmarks or previous analysis. Such a prior helps to stabilize the model and regularize the results toward a reasonable expectation. One challenge that MMMs often face is high variance in model estimates caused by highly correlated input variables (Chan and Perry [2017]). The proposed calibration method provides a means of incorporating ROAS priors which can potentially help to reduce the variance, preventing radical changes in the model estimates due to small changes in the data or the addition or subtraction of seemingly unrelated variables in the model.

The remainder of this paper is organized as follows. In Section 2, we briefly review the GBHMMM and the current practice. In Section 3, we introduce the calibration methodology. Section 3.1 defines the incremental sales estimand. Section 3.2 discusses the calculation of ROAS. Section 3.3 describes the calibration methodology through model reparameterization. Section 3.4 highlights the implications and significance of the proposed methodology. Section 3.5 discusses the adaptability and scalability of the methodology. In Section 4, we evaluate the calibration results through simulation studies. We conclude this paper with a brief discussion in Section 5.

2 Problem Formulation

In this section, we first provide an overview of Geo-level Bayesian Hierarchical Media Mix Modeling (GBHMMM, Sun et al. [2017]), which serves as the foundation for the development of the calibration methodology. Then, we discuss the challenges associated with MMMs calibration within the current practices.

2.1 GBHMMM

For geo g , $g = 1, 2, \dots, G$ at time t , $t = 1, 2, \dots, T$, let $y_{t,g}$ denote the geo-level response variable, $x_{t,m,g}$ denote the media variables for media channel m , $m = 1, 2, \dots, M$, and $z_{t,c,g}$ denote the control variables, $c = 1, 2, \dots, C$. The time-series of media variable is denoted by $x_{t,m,g}^* = \{x_{s,m,g}, s \leq t\}$. The response variable typically represents a KPI (e.g. revenue, online inquires, etc). The media variables could be advertising spend or number of impressions delivered. The control variables encompass a wide range of factors, such as product price, promotions, and macroeconomic factors. Including correct control variables in the model is critical since it plays an essential role in generating unbiased estimates as discussed in section 3.1. The sales and media variables can be scaled by the geo population or target market size. Any control variable that roughly scales with population or market size can also be adjusted to a "per capita" scale. The GBHMMM is modeled as,

$$y_{t,g} = \tau_g + \sum_{m=1}^M \beta_{m,g} \text{Hill}(\text{Adstock}(x_{t,m,g}^*, K_m, S_m), \alpha_m, L) + \sum_{c=1}^C \gamma_{c,g} z_{t,c,g} + \epsilon_{t,g}. \quad (1)$$

The model parameters follow a Bayesian hierarchical structure where each geo is a sample from the overall population and is allowed to deviate from the population level,

$$\beta_{m,g} \stackrel{\text{iid}}{\sim} \text{normal}(\beta_m, \eta_m^2), m = 1, \dots, M,$$

$$\gamma_{c,g} \stackrel{\text{iid}}{\sim} \text{normal}(\gamma_c, \xi_c^2), c = 1, \dots, C,$$

$$\tau_g \stackrel{\text{iid}}{\sim} \text{normal}(\tau, \kappa^2), \epsilon_{t,g} \stackrel{\text{iid}}{\sim} \text{normal}(0, \sigma^2),$$

where β_m , γ_c and τ are hyperparameters that represents the common mechanism of media impact at the total population level. The geo-level variation is controlled by the standard deviations η_m , ξ_c and κ , respectively. Priors are needed for the hyperparameters β_m, γ_c, τ and standard deviations η_m, ξ_c, κ . The shape and carryover effect of advertising is modeled through the Hill function and the geometric Adstock function, and are defined as

$$\text{Hill}(x; K, S) = \frac{1}{1 + (x/K)^{-S}}, \quad (2)$$

$$Adstock(x_0, \dots, x_t; \alpha, L) = \frac{\sum_{l=0}^L \alpha^l x_{t-l}}{\sum_{l=0}^L \alpha^l}, \quad (3)$$

where $\alpha \in (0, 1)$ is the retention rate of the ad effect of the media. The integer L is the maximum duration of carry effect. Hill function is applied after the Adstock transformation to capture the diminishing return of media spend with parameters $K > 0$ and $S > 0$. K is also referred to as EC_{50} , the half saturation point as $Hill(K; K, S) = 1/2$ for any value of K and S . The Hill function goes to 1 as the media spend goes to infinity.

2.2 Current Practice

MMMs analyze aggregated historical data to measure and compare effectiveness across a variety of media channels over the modeling window. Effective model calibration is a critical and indispensable component of the MMM model-building process (Runge et al. [2023]). Calibration involves refining MMMs by incorporating robust measurements, such as incrementality experiments, to ensure that the models provide an accurate and faithful representation of ad effectiveness. Additionally, calibration helps to regularize the model by reducing the uncertainty in the model estimates. To achieve these goal, advertisers often need to design and implement relevant experiments or randomized controlled trials that generate reliable calibration data, such as conversion lift experiments or geo-based lift experiments (Chen et al. [2021]).

Here, we introduce several commonly used methodologies for incorporating relevant experimental results into the calibration process of MMMs. Unregularized non-Bayesian MMMs only allow for calibration after the fact, and use incrementality experiment results to choose the best version of MMMs. During the modeling process, modelers typically need to make important decisions about the assumptions, such as how to represent carryover and shape effects in their model. The assumptions chosen will impact the results of the model. When faced with different models of comparable statistical quality, one can check the proximity of the MMM ROAS and the incrementality experiment iROAS. The least similar results can be used to rule out some models. However, this approach cannot ensure a close alignment between the model’s outputs and the experiment results, and the selection of the final model is subjective and potentially arbitrary.

Modern MMMs often employ Bayesian modeling methodologies to accommodate modelers’ domain knowledge about the effectiveness of various media channels. This domain knowledge is captured by the priors of the model parameters, which are then combined with the information in the raw data to produce the posterior estimation. Incrementality experiments provide excellent opportunities to use priors and enhance the model’s accuracy. However, the process of integrating incrementality experiment results, particularly ROAS, into MMMs via prior adjustment remains ambiguous. A conventional approach is an iterative process that involves attempting various combinations of prior distributions for model parameters such that the implied prior distribution of ROAS aligns with the experiments. This approach can be time-consuming and may necessitate significant expertise.

It is essential to acknowledge that MMM and experiments assess disparate aspects of advertising impact. MMM measures the long-term average effectiveness of ads, encompassing the shape and carryover effect. In contrast, incrementality experiments are usually conducted over a brief period to measure the immediate or short-term effectiveness of ads. This can

make it more challenging to calibrate MMMs using an iterative approach. As such, we are motivated to develop a new methodology to calibrate MMM under the Geo-level Bayesian Hierarchical Media Mix Modeling framework.

3 Methods

3.1 Incremental Sales Estimand

We begin by defining the incremental sales estimand within the potential outcomes framework, which is also known as the Rubin Causal Model or the Rubin Causal Framework (Rubin [2005]), a foundational concept in the field of causal inference. At its core, the Potential Outcome Framework is designed to answer causal questions by considering the counterfactual scenarios, which involve comparing the potential outcome if ad channel turned on during the period $[T_0, T_1]$ to the potential outcome if ad channel turned off during the period $[T_0, T_1]$.

Let $X^{a,m}$ represent the time-series of media variables with the m -th media spend multiplied by a constant a during the period $[T_0, T_1]$, for example, $X^{1,m}$ represents the observed time-series of media variables and $X^{0,m}$ represents the time-series of media variables with the m -th media channel turned off during $[T_0, T_1]$. Let $Z_{t,g} = \{z_{t,c,g}, 1 \leq c \leq C\}$ denote the control variables at geo g and time t , and $Sales_{t,g}^{X^{a,m}}$ denote a random variable for the sales at geo g and time t that would occur if media spend were assigned the value $X^{a,m}$. For any time period t , there are two potential sales outcomes:

$Sales_{t,g}^{X^{1,m}}$: potential outcomes if ad channel executed at historical levels during $[T_0, T_1]$
 $Sales_{t,g}^{X^{0,m}}$: potential outcomes if ad channel turned off during $[T_0, T_1]$

The incremental sales for the media spend over $[T_0, T_1]$ is given by

$$\sum_{t,g} (E[Sales_{t,g}^{X^{1,m}} | Z_{t,g}] - E[Sales_{t,g}^{X^{0,m}} | Z_{t,g}]), \quad (4)$$

The above equation is the estimand of interest, which can be estimated from the conditional expectations from a regression model, such as a Media Mix Model, as in

$$\sum_{t,g} (E[Sales_{t,g} | X^{1,m}, Z_{t,g}] - E[Sales_{t,g} | X^{0,m}, Z_{t,g}]), \quad (5)$$

Note that formulas (5) and (4) are equivalent under the conditional exchangeability assumption and the consistency assumption in casual inference (Hernan and Robins [2020]), as follows,

$$\begin{aligned}
& \sum_{t,g} (E[\text{Sales}_{t,g}^{X^{1,m}} | Z_{t,g}] - E[\text{Sales}_{t,g}^{X^{0,m}} | Z_{t,g}]) \\
\text{exchangeability} & \stackrel{=}{=} \sum_{t,g} (E[\text{Sales}_{t,g}^{X^{1,m}} | X^{1,m}, Z_{t,g}] - E[\text{Sales}_{t,g}^{X^{0,m}} | X^{0,m}, Z_{t,g}]) \\
\text{consistency} & \stackrel{=}{=} \sum_{t,g} (E[\text{Sales}_{t,g} | X^{1,m}, Z_{t,g}] - E[\text{Sales}_{t,g} | X^{0,m}, Z_{t,g}]),
\end{aligned}$$

Conditional exchangeability and consistency of potential outcomes provide the justification for the causal interpretation of a regression model, such as a Media Mix Model. When conditional exchangeability¹ and consistency of potential outcomes² hold, we can obtain unbiased estimates of the causal effect of treatment $X^{a,m}$ on the outcome $\text{Sales}_{t,g}^{X^{a,m}}$ via regression models. Consistency will often hold in the context of a MMM, as the treatment variables are well-defined. Conditional exchangeability will hold if we appropriately control for all confounders³ $Z_{t,g}$ in MMMs. However, it's important to be cautious when selecting or identifying confounders, since in certain circumstances, conditioning on non-confounders (i.e. mediators and colliders according to Pearl [2009]) can introduce bias into the estimates. Moreover, failing to account for confounders that affect both outcome and treatment also introduces bias⁴ into the estimates. As such, subject-matter knowledge becomes necessary to identify possible confounders given that the causal interpretation in MMMs relies on the uncheckable assumption of conditional exchangeability.

3.2 Calculation of ROAS

ROAS is the incremental sales per dollar spent on the media. The ROAS for the media spend over $[T_0, T_1]$ is calculated by dividing the incremental sales estimand as described in equation (4) by the media spend over the period $[T_0, T_1]$. The ROAS for m -th media is defined as

$$\text{ROAS}_m = \frac{\sum_{t,g} (E[\text{Sales}_{t,g}^{X^{1,m}} | Z_{t,g}] - E[\text{Sales}_{t,g}^{X^{0,m}} | Z_{t,g}])}{\sum_{T_0 \leq t \leq T_1} \sum_g C_{t,m,g}}, \quad (6)$$

Where $C_{t,m,g}$ is the media spend for the m th media channel at time t and geo g . Substitute the numerator with equation (5) while accounting for the fact that $E[\text{Sales}_{t,g} | X^{1,m}, Z_{t,g}] - E[\text{Sales}_{t,g} | X^{0,m}, Z_{t,g}]$ for t outside of the range $[T_0, T_1 + L]$ are zero⁵, then

¹Definition of conditional exchangeability: $(\text{Sales}_{t,g}^{X^{1,m}}, \text{Sales}_{t,g}^{X^{0,m}}) \perp X^{a,m} | Z_{t,g}$. It holds in observational studies when the probability of receiving treatment (i.e., media turned on/off) is not dependent on any other unmeasured causes of treatment and outcome, conditional on the measured covariates.

²Definition of consistency: $\text{Sales}_{g,t}^{X^{a,m}} = \text{Sales}_{g,t}$ when $X^{a,m}$ is the observed treatment. It implies that the outcome observed is the exact outcome that was expected.

³In causal inference, a confounder is a variable that influences both the dependent variable and independent variable, causing a spurious association.

⁴Section 4.3 provides more discussion on the bias that arises when exchangeability doesn't hold.

⁵Because the media variables are only changed during the period $[T_0, T_1]$, the difference in the expected sales is non-zero in the range $[T_0, T_1 + L]$ to account for the carryover effect.

$$ROAS_m = \frac{\sum_{T_0 \leq t \leq T_1+L} \sum_g (E[Sales_{t,g}|X^{1,m}, Z_{t,g}] - E[Sales_{t,g}|X^{0,m}, Z_{t,g}])}{\sum_{T_0 \leq t \leq T_1} \sum_g C_{t,m,g}}, \quad (7)$$

The conditional expected sales can be rewritten as the sum of the first three terms in (1), excluding the noise term. If a transformation of sales, such as standardization, is used as the response variable, a corresponding inverse transformation should be applied to the predicted sales⁶. According to the model specification in equation (1), ROAS can be written as

$$ROAS_m = \frac{\sum_{T_0 \leq t \leq T_1+L} \sum_g \beta_{m,g} [Hill(Adstock(x_{t,m,g}^*, \alpha_m, L), K_m, S_m) - Hill(Adstock(\tilde{x}_{t,m,g}^*, \alpha_m, L), K_m, S_m)]}{\sum_{T_0 \leq t \leq T_1} \sum_g C_{t,m,g}} \quad (8)$$

Where $\tilde{x}_{t,m,g}^*$ represents the time-series of m -th media variable at geo g up to time t with the channel turned off during $[T_0, T_1]$. For simplicity, we rewrite the terms in the square brackets as $F(K_m, S_m, \alpha_m)$ since $x_{t,m,g}^*$ and $\tilde{x}_{t,m,g}^*$ are known and L is in general a pre-defined quantity, then,

$$ROAS_m = \frac{\sum_{T_0 \leq t \leq T_1+L} \sum_g \beta_{m,g} F(K_m, S_m, \alpha_m)}{\sum_{T_0 \leq t \leq T_1} \sum_g C_{t,m,g}}. \quad (9)$$

3.3 Model Reparameterization and Calibration

Traditionally, MMM is calibrated by tuning the prior distribution of β_m such that the prior distribution of the ROAS estimate in Equation (9) aligns with the experimental results or prior knowledge. This approach suffers from two drawbacks: First, β_m is assumed to be independent of other parameters, and therefore must be tuned analytically through an iterative process, which is time-consuming and requires significant expertise. Second, the independent prior distribution of β_m may not exist when the desired prior distribution of ROAS has a small standard deviation because β_m is not the only parameter that contributes to the ROAS prior variance. To address the increasing need for better alignment between MMM and experiments, we propose a new MMM calibration approach. The key idea behind this approach is to reparameterize GBHMMM by incorporating ROAS directly as a model parameter.

First, we parameterize $\beta_{m,g}$ as $\beta_{m,g} = \beta_m + \eta_m Z_{m,g}$, where $Z_{m,g} \stackrel{iid}{\sim} normal(0,1)$ and $\{Z_{m,g}\}_{g=1}^G$ is independent of β_m, η_m and all other model parameters. We can then rewrite Equation (9) as

$$ROAS_m = \frac{\sum_{T_0 \leq t \leq T_1+L} \sum_g (\beta_m + \eta_m Z_{m,g}) F(K_m, S_m, \alpha_m)}{\sum_{T_0 \leq t \leq T_1} \sum_g C_{t,m,g}}, \quad (10)$$

⁶For simplicity, equation (8) is assuming that $y_{t,g} = Sales_{t,g}$. In practice, we usually apply a transformation of sales, $y_{t,g} = Lt,g(Sales_{t,g})$, where $Lt,g(\cdot)$ is a linear transformation that includes population scaling and standardization. If a transformation is applied, Equation (8) will need to include an inverse transformation, $Lt,g^{-1}(\cdot)$, in the numerator.

Algebraically, we can write β_m as

$$\begin{aligned}\beta_m &= \frac{\sum_{T_0 \leq t \leq T_1} \sum_g C_{t,m,g} ROAS_m - \sum_{T_0 \leq t \leq T_1+L} \sum_g \eta_m Z_{m,g} F(K_m, S_m, \alpha_m)}{\sum_{T_0 \leq t \leq T_1+L} \sum_g F(K_m, S_m, \alpha_m)} \\ &:= H(ROAS_m, \eta_m, K_m, S_m, \alpha_m, \{Z_{m,g}\}_{g=1}^G),\end{aligned}\tag{11}$$

Given that the historical media spend, $C_{t,m,g}$, is known, and T_0 and T_1 are pre-defined quantities ⁷, β_m is a closed-form function of $\{ROAS_m, \eta_m, K_m, S_m, \alpha_m, \{Z_{m,g}\}_{g=1}^G\}$. For simplicity, we write the closed-form function as $H(ROAS_m, \eta_m, K_m, S_m, \alpha_m, \{Z_{m,g}\}_{g=1}^G)$, which is equal to the first line in (11). Finally, we can reparameterize the GBHMMM in (1) with $ROAS_m$ as a parameter instead of β_m

$$\begin{aligned}y_{t,g} &= \sum_{m=1}^M (H(ROAS_m, \eta_m, K_m, S_m, \alpha_m, \{Z_{m,g}\}_{g=1}^G) + \eta_m Z_{m,g}) Hill(Adstock(x_{t,m,g}^*, \alpha_m, L), K_m, S_m) \\ &\quad + \tau_g + \sum_{c=1}^C \gamma_{c,g} z_{t,c,g} + \epsilon_{t,g}.\end{aligned}\tag{12}$$

In contrast to the original GBHMMM which places priors on β_m , the proposed calibration places priors on $ROAS_m$. To a finer point, the calibration implicitly places a prior on $\beta_m = H(ROAS_m, \eta_m, K_m, S_m, \alpha_m, \{Z_{m,g}\}_{g=1}^G)$, which is determined by the priors selected for $\{ROAS_m, \eta_m, K_m, S_m, \alpha_m, \{Z_{m,g}\}_{g=1}^G\}$. In particular, the prior distributions of β_m and $\{Z_{m,g}\}_{g=1}^G$ are no longer independent. In the reparameterized model, we now define $Z_{m,g} \stackrel{\text{iid}}{\sim} normal(0, 1)$ and $\{Z_{m,g}\}_{g=1}^G$ is independent of $ROAS_m, \eta_m$ and all other model parameters.

3.4 Implications and significance

Calibrating the MMM with ROAS priors is appealing in practice because it provides a straightforward way to incorporate the prior domain knowledge about the effectiveness of media channels and the incrementality experiment results into MMM. It would directly answer many of the fundamental questions advertisers have: How do I incorporate the results of a single incrementality experiment into MMMs? How do I calibrate MMMs when the experiment captures only short-term effects and doesn't represent modeling-window ROAS accurately? How should MMM calibration be approached when multiple incrementality experiments have been conducted for a single media channel?

Incorporate the results of an incrementality experiment

Incrementality experiments, such as Google trimmed match geo experiments (Chen et al. [2021]), typically indicate a range of plausible values (e.g., mean, confidence interval, p-values) that can inform our choice of prior distribution for $ROAS_m$ in Equation (12), provided that the estimand of the experiment aligns with the definition in (4).

⁷The selection of T_0 and T_1 is further discussed in Section 3.4.

When the experiment captures only short-term effects

However, it's worth noting that the experimental results may not reflect the modeling-window ROAS that MMM aims to measure, particularly when the media spend exhibits strong seasonality and the incrementality experiments have a short measurement window. In this case, the experiment is assessing a specific point on the saturation curve that may not be representative of overall media performance. Figure 1 illustrates that for a given media channel, a typical C-shaped saturation or diminishing return curve would yield a higher ROAS for lower-than-usual spend during the experiment window, as compared to the MMM modeling time window. The calibration method offers the flexibility for users to calibrate the MMM using a subset of data by setting $[T_0, T_1]$ as the experiment period, or to calibrate the MMM using the complete data by setting $T_0 = 0$ and $T_1 = T$, depending on whether the experimental results accurately represent the average media effect within the modeling window.




Figure 1: Isolating experiment-window ROAS from diminishing return

When multiple incrementality experiments are available

When multiple incrementality experiments have been run for a single media, an intuitive way to calibrate GBHMMM is to amend β_m in equation (11) as

$$\beta_m = \frac{\sum_e \sum_{T_{e,0} \leq t \leq T_{e,1}} \sum_g C_{t,m,g} ROAS_{e,m} - \sum_e \sum_{T_{e,0} \leq t \leq T_{e,1}+L} \sum_g \eta_m Z_{m,g} F(K_m, S_m, \alpha_m)}{\sum_e \sum_{T_{e,0} \leq t \leq T_{e,1}+L} \sum_g F(K_m, S_m, \alpha_m)},$$

where $ROAS_{e,m}$ is ROAS parameter for the e -th experiment run for media m , and $[T_{e,0}, T_{e,1}]$ are the start and end dates for the e -th experiment. However, we have noticed that the parameters $\{ROAS_{e,m}\}_{e=1}^E$ are essentially unidentifiable in some scenarios. For example, different sets of $\{ROAS_{e,m}\}_{e=1}^E$ can lead to same value of β_m . The poor identifiability of $\{ROAS_{e,m}\}_{e=1}^E$ makes it challenging to estimate the parameters well with any statistical method. Therefore, we suggest replacing $\{ROAS_{e,m}\}_{e=1}^E$ with \overline{ROAS}_m , which represents the aggregated ROAS for media m during all experiment periods, then,

$$\beta_m = \frac{\sum_e \sum_{T_{e,0} \leq t \leq T_{e,1}} \sum_g C_{t,m,g} \overline{ROAS}_m - \sum_e \sum_{T_{e,0} \leq t \leq T_{e,1}+L} \sum_g \eta_m Z_{m,g} F(K_m, S_m, \alpha_m)}{\sum_e \sum_{T_{e,0} \leq t \leq T_{e,1}+L} \sum_g F(K_m, S_m, \alpha_m)}, \quad (13)$$

One way to obtain the prior distribution of \overline{ROAS}_m when multiple experiments exist is to take the weighted sum of the prior ROAS distributions for each experiment, where the weights, denoted by $w_{e,m}$, are proportional to the media spend during the experiments, as in

$$\begin{aligned} Pr(\overline{ROAS}_m \leq R) &= Pr(\sum_e w_{e,m} ROAS_{e,m} \leq R), \\ w_{e,m} &= \frac{\sum_{T_{e,0} \leq t \leq T_{e,1}} \sum_g C_{t,m,g}}{\sum_e \sum_{T_{e,0} \leq t \leq T_{e,1}} \sum_g C_{t,m,g}} \end{aligned}$$

3.5 Adaptability and scalability

The proposed calibration approach is not limited to the GBHMMM, but is also applicable to other Bayesian-based MMM methodologies, such as the BMMM (Jin et al. [2017]) and RFMMM (Zhang et al. [2023]), as well as other model specifications, such as the GBHMMM with log-normal $\beta_{m,g}$ prior.

BMMM (Jin et al. [2017]) is modeled by the following generic equation,

$$y_t = \tau + \sum_{m=1}^M \beta_m Hill(Adstock(x_{t,m}^*, \alpha_m, L), K_m, S_m) + \sum_{c=1}^C \gamma_c z_{t,c} + \epsilon_t \quad (14)$$

Following the same reparameterization process, β_m can be written as,

$$\begin{aligned} \beta_m &= \frac{\sum_{T_0 \leq t \leq T_1} C_{t,m} ROAS_m}{\sum_{T_0 \leq t \leq T_1+L} (Hill(Adstock(x_{t,m}^*, \alpha_m, L), K_m, S_m) - Hill(Adstock(\tilde{x}_{t,m}^*, \alpha_m, L), K_m, S_m))} \\ &:= H'(ROAS_m, K_m, S_m, \alpha_m), \end{aligned}$$

As a result, BMMM can also be reparameterized with $ROAS_m$ as a parameter instead of β_m , as in

$$y_t = \tau + \sum_{m=1}^M H'(ROAS_m, K_m, S_m, \alpha_m) Hill(Adstock(x_{t,m}^*, \alpha_m, L), K_m, S_m) + \sum_{c=1}^C \gamma_c z_{t,c} + \epsilon_t \quad (15)$$

GBHMMM with log-normal prior on $\beta_{m,g} \stackrel{iid}{\sim} \text{lognormal}(\beta_m, \eta_m^2)$ is in general used to prevent negative media effects, where $\beta_{m,g}$ can be parameterized as $\beta_{m,g} = e^{\beta_m + \eta_m Z_{m,g}}$. Following the same reparameterization process, β_m can be written as,

$$\begin{aligned}\beta_m &= \log \left(\sum_{T_0 \leq t \leq T_1} \sum_g C_{t,m,g} ROAS_m \right) - \log \left(\sum_{T_0 \leq t \leq T_1+L} \sum_g e^{\eta_m Z_{m,g}} F(K_m, S_m, \alpha_m) \right) \\ &:= H''(ROAS_m, \eta_m, K_m, S_m, \alpha_m, \{Z_{m,g}\}_{g=1}^G),\end{aligned}$$

where $F(K_m, S_m, \alpha_m)$ is the same function as defined in Equation (9). Then, GBHMMM with log-normal $\beta_{m,g}$ prior can also be reparameterized with $ROAS_m$ as a parameter instead of β_m , as in

$$\begin{aligned}y_{t,g} &= \sum_{m=1}^M e^{H''(ROAS_m, \eta_m, K_m, S_m, \alpha_m, \{Z_{m,g}\}_{g=1}^G) + \eta_m Z_{m,g}} Hill(Adstock(x_{t,m,g}^*, \alpha_m, L), K_m, S_m) \\ &\quad + \tau_g + \sum_{c=1}^C \gamma_{c,g} z_{t,c,g} + \epsilon_{t,g}.\end{aligned}\tag{16}$$

In summary, the presented reparameterization and calibration approach can be applied to other Bayesian based MMMs. This approach is viable for implementation as long as there is a one-to-one mapping between ROAS and one of the model parameters, holding other model parameters fixed.

4 Simulation

This section illustrates some of the key properties of the calibration method using simulated data generated from the same model class as in Equation (1). Specifically, we compare the ROAS estimates with and without calibration, and show that the calibrated MMM generally has tighter credible intervals and lower bias for the average ROAS than the non-calibrated MMM.

4.1 Data simulation

We assume that there is only one media variable ($M = 1$), one control variable ($C = 1$), 100 geos ($G = 100$) and two years (104 Weeks) of data. The control variable is simulated first to account for its potential influence as a confounding factor on both media spending and sales. In general, control variables may exhibit seasonal patterns, so we simulated the control variable $z_{t,g}$ as a sinusoid over time with addition of random noise, $z_{t,g} = \cos(2\pi(t - 12)/52) + w_{t,g}$, where $w_{t,g}$ is white noise with $w_{t,g} \stackrel{\text{iid}}{\sim} normal(0, 1)$.

We then simulate the media spend per capita in the geo g to have a positive correlation ρ_g with the control variable,

$$\begin{aligned}x_{t,g} &= u_g + \rho_g z_{t,g} + \sqrt{1 - \rho_g^2} v_{t,g}, \\ u_g &\stackrel{\text{iid}}{\sim} normal(0, 1), \rho_g \stackrel{\text{iid}}{\sim} uniform(0, 1), v_{t,g} \stackrel{\text{iid}}{\sim} normal(0, 1)\end{aligned}$$

As the media spend should be non-negative, we take the positive part of $x_{t,g}$ as the media spend. Finally, the geo level sales per capita $y_{t,g}$ is simulated to depend on the control variable and the media spend, as in

$$y_{t,g} = \tau_g + \beta_g \text{Adstock}(\text{Hill}^*(x_{t,g}^*, K, S), \alpha, L) + \gamma_g z_{t,g} + \epsilon_{t,g}$$

$$\tau_g \stackrel{\text{iid}}{\sim} \text{normal}(\tau, \kappa^2), \beta_g \stackrel{\text{iid}}{\sim} \text{normal}(\beta, \eta^2), \gamma_g \stackrel{\text{iid}}{\sim} \text{normal}(\gamma, \xi^2), \epsilon_{t,g} \stackrel{\text{iid}}{\sim} \text{normal}(0, \sigma^2)$$

where the hyper-parameters are fixed and summarized in Table 1, and the remaining parameters $(\tau_g, \beta_g, \gamma_g, \epsilon_{t,g})$ are randomly drawn from the above distributions.

α	K	S	L
0.5	1	1	3
τ	β	γ	
10	3	1	
κ	η	ξ	σ
1	1	1	5

Table 1: Hyper-parameters

The ground-truth ROAS for the simulated data over a specified time period can be calculated by applying Equation (8) with the parameter values obtained from the simulation. We partition the two-year dataset into eight quarters and illustrate the ground-truth ROAS for each quarter, along with the time series plots of the sales and the media spends for one of the geos, as presented in Figure 2.

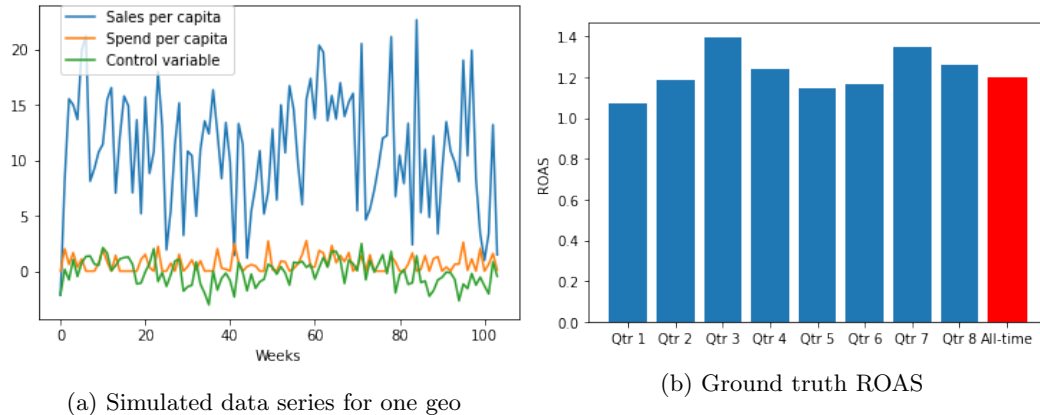



Figure 2: Simulated data and ground truth ROAS

4.2 ROAS estimates with and without calibration

In this section, we compare the results from the GBHMMM (Equation 1) using non-informative priors with those from the calibrated MMM (Equation 12) using ROAS prior. As described in Section 3.4, the MMM can be calibrated through two approaches: either using the ROAS prior specific to a sub-window or employing the ROAS prior associated with the entire modeling window. For comparison, we obtain both the ROAS prior for the entire

modeling window and the ROAS prior for the sub-window of quarter 3 by repeating the data simulation process 100 times, and then matching the mean and the standard deviation of the simulated ROAS to a log-normal distribution ⁸, as illustrated in Figure 3.



(a) ROAS Prior for the entire modeling window (b) ROAS Prior for the sub-window of quarter 3

Figure 3: Simulated data and ground truth ROAS

Figure 4 displays the posterior ROAS estimates with and without calibration. The blue curve is the posterior density of ROAS for GBHMMM without incorporating ROAS priors, the orange curve is the posterior density of ROAS when the sub-window ROAS prior is applied, the green curve is the posterior density of ROAS when using the full-window ROAS prior, and the red line is the ground-truth ROAS of the entire modeling window. It is observed that the GBHMMM without calibration yields ROAS estimates with greater uncertainty in comparison to the GBHMMM calibrated using ROAS priors. Specifically, the standard deviation of the posterior ROAS estimates from the GBHMMM without calibration is approximately five times that of the GBHMMM with calibration. In contrast, when calibrating the GBHMMM using either the ROAS prior for the entire modeling window or the ROAS prior for the sub-window of quarter 3, the posterior ROAS distributions exhibit a substantial degree of similarity, with the full-window calibration yielding a slightly narrower credible interval.

This study is an ideal case in the sense that it is free of omitted variables and model mis-specification, which often exist in real data. In the following section, we will test the calibration result in more complex settings mimicking the challenges faced by modelers in practice.

4.3 When bias exists

In practice, when the confounding factor, $z_{t,g}$, is unobserved, it's common to directly adjust for the available proxies in order to reduce bias due to unobserved confounding. However, the bias could not be entirely eliminated when the confounding factor is unobserved, and can only be measured with error via a proxy variable (Kurok and Pea [2014]). To mimic this situation, we simulate a proxy control variable, $z_{t,g}^{proxy} = z_{t,g} + u_{t,g}$, where $u_{t,g}$ is white noise independent of $z_{t,g}$ and $u_{t,g} \stackrel{iid}{\sim} normal(0,1)$. Assuming $z_{t,g}$ is unobserved, we use $z_{t,g}^{proxy}$

⁸Lognormal distribution is in general used for setting ROAS prior to prevent negative ROAS estimates. It is important to ensure that the ROAS prior distribution is centered around the ground truth ROAS in Figure 2, which might not always be the case as the ground truth ROAS is from a single simulation replicate.

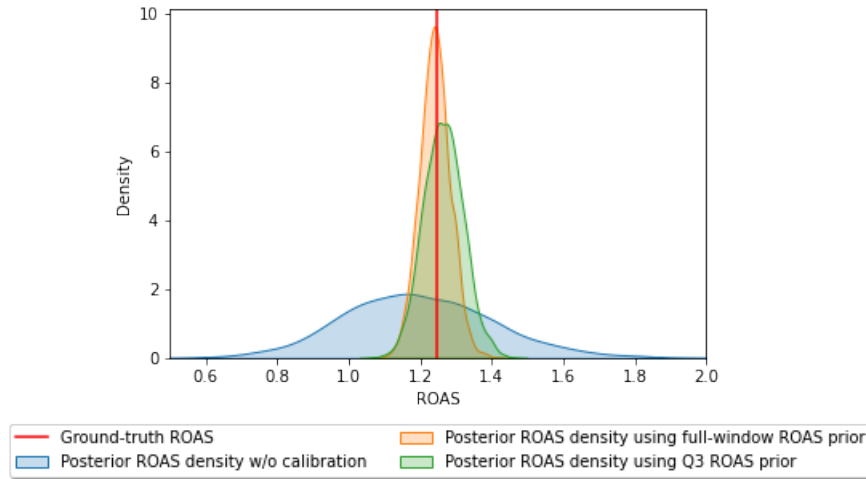


Figure 4: Posterior ROAS vs ground truth

instead of $z_{t,g}$ to construct models and evaluate the impact of calibration using ROAS priors.

Figure 5 shows that the GBHMM, whether calibrated or not, produces biased estimates of ROAS when the confounding factor is unobserved and a proxy variable is used. However, the use of ROAS priors for calibration proves to significantly reduce this bias. It is also worth noting that calibration using the ROAS prior for the entire modeling window is shown to be more effective in reducing bias when compared to calibration with sub-window ROAS priors.

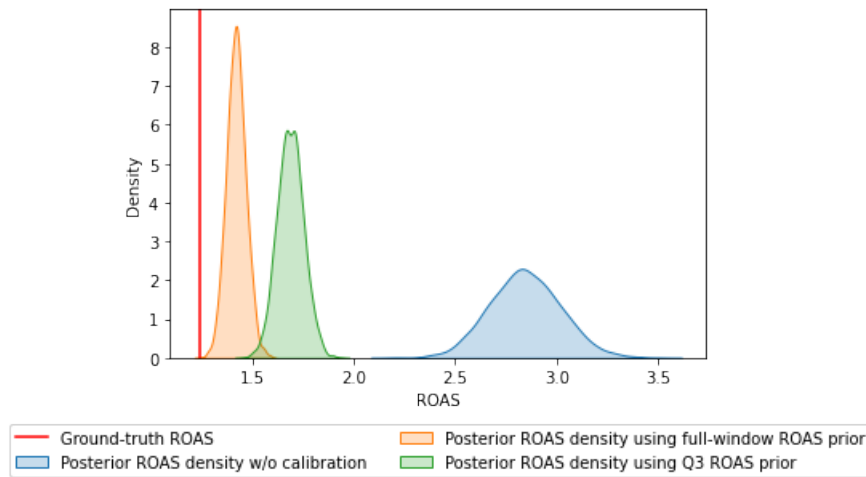


Figure 5: Posterior ROAS vs ground truth

5 Conclusion

Media Mix Models (MMMs) frequently integrate Bayesian modeling methodologies due to their flexibility in accommodating prior domain knowledge regarding the effectiveness of media channels. Until now, most models do not include ROAS as a parameter, which makes it difficult to translate prior information about ROAS into an appropriate prior distribution. In this paper, we propose a Bayesian MMM that includes ROAS as a model parameter. This provides a straightforward way to calibrate MMMs based on any available prior information, including domain knowledge or incrementality experiments. We also discuss how the MMM calibration should be approached when the incrementality experiment measures only short-term ad effect and doesn't represent the modeling-window ROAS, and when multiple incrementality experiments are conducted for a single media channel. The proposed calibration method can be generalized to other Bayesian-based MMMs as long as one of the model parameters can be parameterized as a function of ROAS.

The simulation study in section 4 provides evidence that, by employing ROAS priors through reparameterization for MMM calibration, this methodology can significantly reduce the bias and uncertainty in the resultant posterior ROAS estimates. Furthermore, the study demonstrates that calibration with modeling-window ROAS priors is more effective in reducing bias when compared to calibration with short-term ROAS priors.

Acknowledgment

We would like to thank Hitesh Wadhvani and Lisa Jakobovits for their helpful comments and suggestions, Dirk Nachbar, Saket Kumar and Google gTech Ads colleagues for their great discussion and feedback, Harikesh Nair for his guidance and support throughout the research, Vivek Beniwal, Yonathan Schwarzkopf and Eng colleagues for their help in implementing the method.

References

- David Chan and Mike Perry. Challenges and opportunities in media mix modeling. *research.google*, 2017.
- Aiyue Chen, Marco Longfils, and Nicolas Remy. Trimmed match design for randomized paired geo experiments. Technical report, Google LLC, 2021. URL <https://arxiv.org/abs/2105.07060>.
- Miguel A. Hernan and James Robins. *Causal Inference: What If*. Chapman & Hall/CRC, 2020. ISBN 9781420076165.
- Yuxue Jin, Yueqing Wang, Yunting Sun, David Chan, and Jim Koehler. Bayesian methods for media mix modeling with carryover and shape effects. *research.google*, 2017.
- Manabu Kurok and Judea Pea. Measurement bias and effect restoration in causal inference. *Biometrika*, 2014.
- Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2009. ISBN 9780521895606.

- Donald B. Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 2005.
- Julian Runge, Harpreet Patter, and Igor Skokan. A new gold standard for digital ad measurement? <https://hbr.org/2023/03/a-new-gold-standard-for-digital-ad-measurement>, 2023.
- Yunting Sun, Yueqing Wang, Yuxue Jin, David Chan, and Jim Koehler. Geo-level bayesian hierarchical media mix modeling. *research.google*, 2017.
- Yingxiang Zhang, Mike Wurm, Alexander Wakim, Eddie Li, and Ying Liu. Bayesian hierarchical media mix model incorporating reach and frequency data. *research.google*, 2023.